

(43) Date of A Publication 14.10.1998

(21) Application No 9707316.7

(22) Date of Filing 10.04.1997

(71) Applicant(s)
International Business Machines Corporation
(Incorporated in USA - New York)
Armonk, New York 10504, United States of America

(72) Inventor(s)
Adrian Mark Colyer
Andrew Key
Vincent Sethi

(74) Agent and/or Address for Service
P Waldner
IBM United Kingdom Limited, Intellectual Property
Department, Hursley Park, WINCHESTER, Hampshire,
SO21 2JN, United Kingdom

(51) INT CL⁶
G06F 9/46 17/60

(52) UK CL (Edition P)
G4A APX AUXX

(56) Documents Cited
GB 2281423 A EP 0657833 A2 EP 0645726 A2
EP 0645725 A2 EP 0558006 A2 EP 0319232 A2
WO 96/12241 A1 WO 94/11806 A1 US 5392400 A
US 5319777 A US 5293619 A

(58) Field of Search
UK CL (Edition O) G4A APX AUXX
INT CL⁶ G06F 9/46 17/60
Online: WPI

(54) Abstract Title
Personal conferencing system

(57) There is described a personal conferencing method and system for a client (32)/server (30) environment. The server (30) stores conference model data (60) such as a shared chalkboard or a molecular model and each client has a copy of the conference model data (60). When one of said clients (32A, 32B, 32C) edits the model (60) it creates an instruction (64D) for operating on the model and sends the instruction (64A) to the server (30). The server (30) operates (62) on its conference model data (60) on receipt of the instruction (64A) and resends the instruction (64B) to each of the clients (32A, 32B, 32C) and each client (32A, 32B, 32C) performs the same operation (62) on their respective copies of the conference data model (60). Whereby after a plurality of different operating instructions (64B) from different clients (32A, 32B, 32C) the respective copies of the conference model data (60) are equivalent.

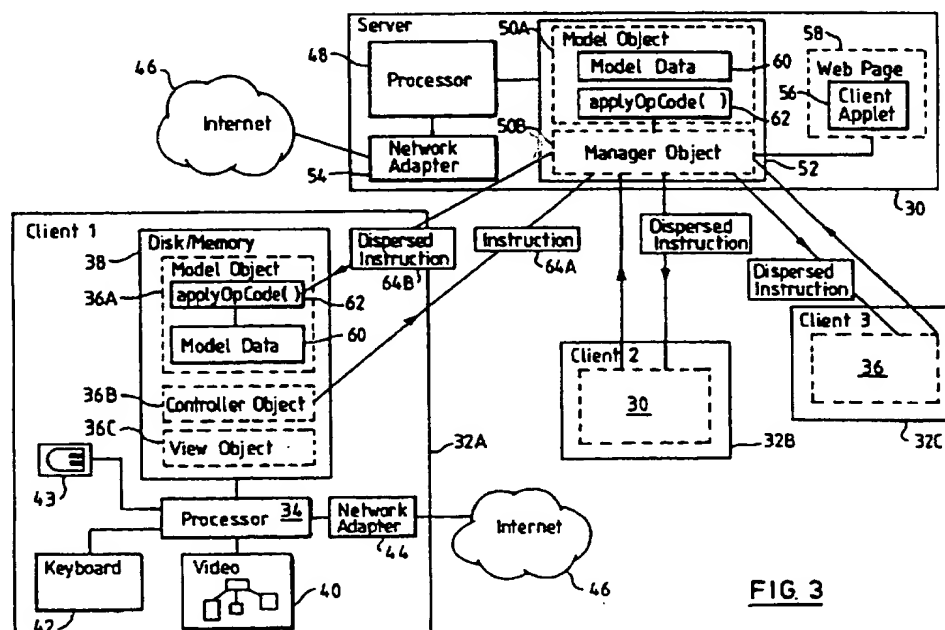
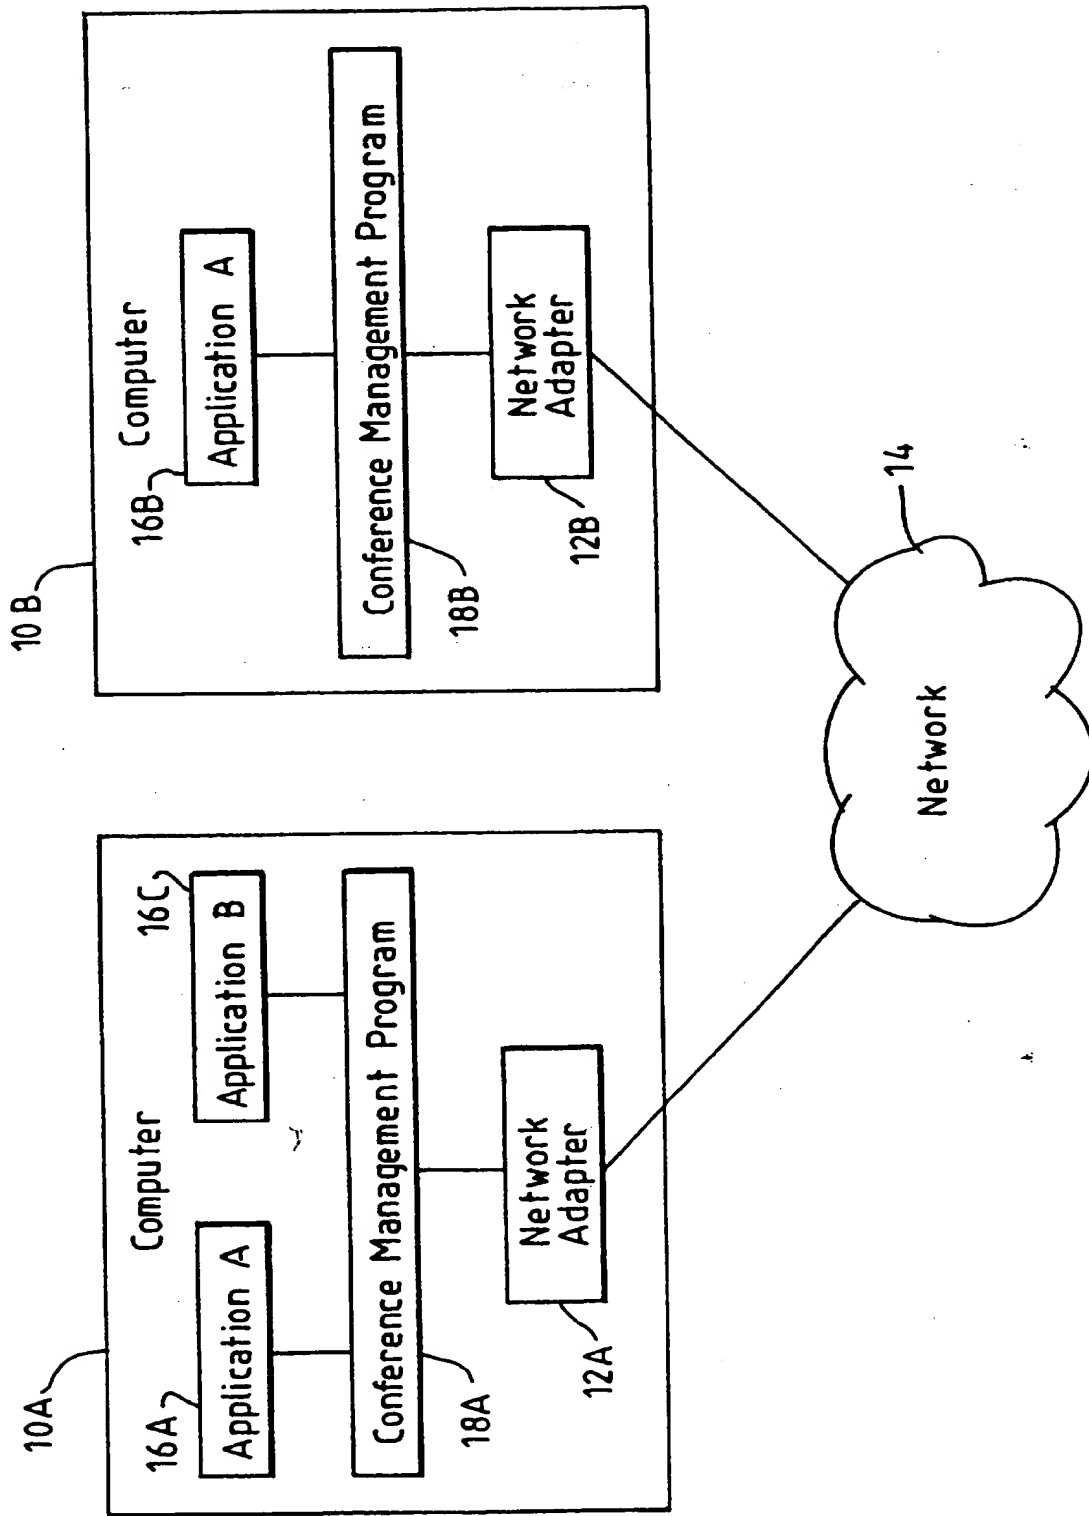


FIG 3

FIG. 1 (PRIOR ART)




FIG. 2A (PRIOR ART)




FIG. 2B (PRIOR ART)




FIG. 3




FIG. 4



FIG. 5

PERSONAL CONFERENCING SYSTEM

This invention relates to a personal conferencing system.

5 Personal computer conferencing allows a computer user to participate in a meeting from his own desk instead of travelling to a location mutually convenient to all parties. Just as telephony systems have developed to include conference calls to more than two parties so have computer systems developed to allow large numbers of networked
10 computer users to participate in a single conference. An example of personal conferencing is an internet 'chat room' in which one can conduct meetings in text or voice in real time with other computer users.

15 A development of the 'chat room' type conference is the application type conference in which Personal Conferencing systems try to allow a number of participants to share applications so that they can all interact with them and see the same results. One example of this is the 'white board' application in which participants can draw on a 'common' white board and see and edit the same drawing. Another example is a
20 molecular modelling application in which each participant may change and manipulate a molecule displayed on the computer screen.

25 Some Personal Conferencing Systems are implemented in a peer-to-peer fashion. An example of this are the PM2PM and Person to Person/2 products developed by IBM Corporation. In this environment, for example (see Figure 1) a conference management/call manager program (18A, 18B) runs in each computer (10A, 10B) drawing upon the services of one or more network adaptors (12A, 12B) to communicate over a network (14). A number of specially written personal conferencing applications (16A, 16B, 16C)
30 connect to the conference management. If a conference includes a shared chalkboard, for example, then on each participating node chalkboard personal conferencing applications (16A, 16B, 16C) will be running. These applications send packets of data to each other, utilising the services of the conference management program (18A, 18B).

35 This arrangement has the following drawbacks:

- (i) As the conferencing management program is not responsible for ensuring that the data in the application on each node corresponding to the same shared chalkboard (or other data)
40 is the same, the applications must invent and use a protocol between them for doing so.

- (ii) Each node may elect to not store an entirely identical copy of the data which its peers do. This design decision is often made in error, as it tends to expedite a rapid implementation, with the side effects not being immediately apparent.
- (iii) If a new participant joins the conference, then a copy of the data in the chalkboard (or other data) must be transmitted from one of the nodes. If the data on each node is not identical, then which data is best sent? For an efficient transfer, the application may need an appreciation of the network geography, further complicating the implementation of the application. When a new participant joins the conference, and wishes to share an existing application he must be brought into step.
- (iv) All this extra complication in the applications (peer-to-peer protocol to keep data in step, joining and leaving transfer, knowledge of geography) must be duplicated in each sort of application.
- (v) A person may not indicate their readiness to participate in a conference on their own, as there is no 'meeting point' inherent in the system. There is no server which can reasonably expect to be powered up at any time, in a purely peer-to-peer system.

An example of personal conferencing which does not allow manipulation of shared data by either party, is simplified Peer-to-Peer such as the shared whiteboard application in the NeWI environment developed by Data Connection Limited. The whiteboard has multiple planes, one per participant. Each participant can draw upon his own plane, and changes to his plane is transmitted to the other participants, who can then see (but not change) the new plane. As nobody can change the picture on the plane of another, there is no shared manipulation of shared data. 'NeWI' is a trademark of Data Connection Limited.

Simplified Peer-to-Peer has further disadvantages. Although in the example of a shared whiteboard some utility is afforded by such a tool, most of the value is in the shared manipulation of data. Although this system is simple to implement, this style of 'conferencing' is awkward for text editing and other forms of shared manipulation.

One of the main problems of peer-to-peer conferencing systems is therefore that of serialisation, the requirement that all the data must arrive at all the destinations/nodes in the same order to preserve the

conference model integrity. The order used is normally the order in which the data is sent.

A different type of personal conferencing is provided by the XTV system developed by IBM Corporation. This is an enhancement/modification of the XWindows system. If the input events (mouse-clicks, keystrokes etc.) from a variety of participants can be transmitted to a given machine, applied to a running program, and the resulting changes to the screen (or windows of the program) transmitted back to the participants, a rudimentary form of personal conferencing can be achieved.

XWindows is a client-server system, whereby the servers serve the ability to provide user input, and also serve the ability to draw to their screen. Typically each user has an Xserver (20) running on their machine and the data-stream is sent to the Xserver (20) via an XTV interceptor (22) to get it to display graphics using X-protocol. As XWindows is client-server, and X-protocol is transmitted between them, it is usual for clients and servers to be running in different machines. XWindows implies networking, and this is handy when collaborative working is desired (as in personal conferencing). For output, XTV (22) intercepts the X-protocol sent by an application running on a client (24) to a server (26). It fans this data out to a number of Xservers (20), thus making the programs windows visible to a number of participants. Similarly, input is drawn from a number of Xservers (20), and channelled in to the machine running the shared (multiply accessible) application. XTV and XWindows are trademarks of IBM Corporation.

Normal XWindows client-server interaction is shown in Figure 2A and Interaction with XTV (22) 'in the middle' is shown in Figure 2B. XTV (22) routes input and output traffic between a number of applications and servers and has the advantage of working with all existing non-collaborative applications.

However, it has the following drawbacks:

- (i) Its not a collaborative application 'enabler'.
- (ii) Non-collaborative applications are written with one user in mind and keystrokes and mouseclicks are from a variety of users, all interleaved in time.
- (iii) All users can only see an identical view of the data. eg:User one cannot see a zoomed view of the chalkboard, whilst user two sees the normal size.

- (iv) Changes to the application data can result in large quantities of screen changes (and thus X-protocol), and this problem is made worse by the fact that it is now sent to all the participants.
- (v) The application is really only running on one machine, and so the data it holds is only accessible on that machine unless all the machines share a common file-system (such as AFS or DFS). Even if machines share access to a common file-system, drive letters and pathnames may be different across machines, for the same file. Generally, when manipulation remote data, you (the user) have to consider that you are running on the remote machine. This is added complexity for the user.

Another problem of most types of Personal Conference system is that no attempt has been made to optimize the network performance by choosing the nearest (network speed wise) existing participant to the new participant. Commonly, personal conference systems have transmitted the data from an arbitrary existing node to a new node.

EPO497022 Hewlett-Packard relates to a distributed object-based computer system in which sharable objects are split into client and server components. It is designed for a number of servers which communicate amongst themselves to send messages and deliver updates.

The embodiment of the invention draws on a number of disciplines for its performance. One of these Model/View/Controller (or MVC for short) is a discipline used in implementing applications. Such an application is divided up into three logical parts, with clearly defined interfaces and responsibilities.

The Model is the data which the application is presenting and allowing to be manipulated. However the Model itself only contains the data and entypoints (ie: methods) by which the data may be changed. The Model has no concept of how the data is to be presented to the user, and no concept of how a user indicates or is allowed to specify their intended change/modification of the data. In the example of a Chalkboard application, the Model could be a bitmap representing the picture. The Model would not include any concept of a current pen colour, fill style or font setting etc. Modification of the data in the Model is achieved through the entypoints/methods, and the new state of the Model is determined entirely from the old state and the parameters to the entypoint (ie: the entypoints are pure, and do not use global data).

In this description the entrypoints/methods are associated with operating on the data of the model or the 'means for operating'.

5 The View is the user-interface components for showing the user the current state of the Model. It determines how to display the model purely from the Model data it is configured to show, with access to no external or global variables. That is to say that given a particular Model, you will always get the same visual representation. In the example of a Chalkboard application, the View might chose to display the
10 Models bitmap in the client area of a window, with scrollbars.

 The Controller is the user-interface components for allowing the user to specify changes to the Model. In the example of a Chalkboard application, the Controller is responsible for showing the rubberbanding
15 which occurs when the user decides to draw a rectangle (or other graphic symbol) on the Chalkboard. The Controller does not directly modify the Model (eg: draw the rectangle) - rather it calls entrypoints in the Model to get this job done. The Controller can maintain the users current choice of pen colour and font, and this information is passed as a part
20 of a draw-rectangle command passed to the Model (in order that the change be completely specified). In this description the controller is associated with the creation of an operating instruction or the means for creating.

25 In the interests of efficiency, when a Model has changed, it may notify the View (or Views) of itself of the fact that it has changed, and perhaps the extent of the change. This allows the View to redraw itself to reflect the change, or even to selectively redraw the part of itself
30 that has been affected.

 In practice the distinction between Controller and View is often blurred. In the case of the Chalkboard, the same client window might both display the current View of the bitmap and the rubberbanding taking
35 place due to the current user interaction with the Controller.

 The embodiments of the present invention attempt to address these and other problems.

40 A personal conferencing method for a server computer (30) having one or more client computers (32A, 32B, 32C) connected comprising the steps of:

- (a) storing conference model data (60) on the server computer (30);
- (b) operating (62) on the conference model data (60) on receipt of an instruction (64A) from a client computer;
- (c) resending the instruction (64B) from the server computer (30) to each of the client computers (32A, 32B, 32C) whereby each client computer (32A, 32B, 32C) performs the same operation (62) on a copy of the conference data model held on the client computer so that all the copies of the conference data model (60) correspond to the conference model data held on the server computer.

Using a client/server configuration for a conferencing system in which each client has a conference model which is updated from the server rather than the client has clear advantages. A serialisation mechanism is provided which is essential for any (non-broken) personal conference system design. All editing messages pass through the server computer, on its way to the client computers, thereby providing a central serialization point. The server computer can promise to deliver data to all client computers in the same order which guarantees that all client computers display the same thing.

Further advantages are apparent. For instance, the design of a personal conferencing system is formalized and simplified through incorporating a Model/View/Controller configuration. Also intercepting changes to the conference model and distributing the change to the client computers is the simplest way of achieving corresponding screen views. In such a design the server computer is free of any user interface code.

The server computer itself provides a permanent 'meeting point', where clients can advertise their existence and willingness to participate in shared work. This advantage is particularly lacking in peer-to-peer systems.

According to a second aspect of the present invention there is provided a personal conferencing method for one or more client computers connected to a server computer comprising the steps of:

- (a) each of said client computers storing a copy of conference model data;
- (b) one of said client computers creating an instruction for operating on the conference model data and sending said instruction to the server computer;

(c) each of said client computers receiving a copy of the instruction from the server computer and operating on their respective models according to the instruction;

whereby after a plurality of different operating instructions from different clients the respective copies of the conference model data correspond with each other.

According to a third aspect of the present invention there is provided a personal conferencing system comprising at least two client computers and a server computer;

the server computer comprising:

conference model data;

means for receiving an instruction from a client computer and operating on the conference model data according to the instruction;

the client computer comprising:

means for creating an instruction and sending the instruction to the server computer;

characterised by:

the server computer further comprising means for dispersing a received instruction to all the client computer; and

each of the client computers further comprising:

a copy of the conference model data;

means for receiving a dispersed instruction from the server computer and operating on the copy of the conference model data according to the instruction received from the server computer.

According to a fourth aspect of the present invention there is provided a personal conferencing system comprising at least two computers connected for message communication; each computer comprising:

conference model data;

means for creating an instruction for operating on the conference model data and sending the instruction to another computer;

means for operating on the conference model data according to an instruction;

characterised by:

one of the computers is a server computer and the other computer or computers are client computers, each having a connection to the server computer;

the server computers further comprises a means for dispersing a received instruction to all the connected client computers; and

the means for operating comprises means for receiving a dispersed instruction from the server computer;

whereby a sequence of operating instructions from different computers in the personal conferencing system will be performed in the same order on each copy of conference model data.

A personal conferencing system comprising:

a server object and a plurality of client objects;

a model object (36A) having model data (60) and means (62) for operating on the model data (60) to an instruction message (64B);

each client object (36) having means (36B) for creating an instruction for operating on the model object (36A);

characterised by:

each client object (36) comprising an associated model object (36A);

each client object (36) further comprises means (36B) for passing an operating instruction to the server object;

said server object (50) further comprises means (50B) for dispersing a passed instruction from the server object to all client objects (36);

thereby allowing a sequence of instructions from different client objects (36) to be performed in the same order by each model object on each instance of the model data (60).

Preferably the instruction is in the form of an Opcode having arguments. This is a convenient way of packaging modifications to the conference model data so that the modifications may be easily distributed from the client computer creating it to the server computer and from the server computer to all the client computers involved in the conference. Clearly if the same initial conference model data exists on a number of client computers, applying the same operation to all the copies of the conference model data will result in the same model.

Preferably the server computer sends a copy of the conference model data to a client computer who wishes to join the conference and makes such a request to the server computer. This enables new client computers to join the conference at any point and to receive an accurate and reliable copy of the conference proceedings at that point. In a preferred embodiment a blank template conference model is stored on the client computer and its state is changed to match the state of the conference model data held in the server.

In the preferred embodiment the methods that are performed in the server computer and the client computers are under the control of a hosted byte code such as JAVA. This allows cross platform portability of a conference. It is advantageous that the machine hosted byte code is downloadable from the server computer by a network browser running on the client computer. This allows easy connectivity by most computers already connected to the internet by real time downloading from the World Wide Web without the need for any pre-installed client code.

In the preferred embodiment each client computer graphically displays the conference model data on a computer screen. Each client computer may therefore adapt the particulars of the way in which the conference model data is displayed on its own screen, i.e. the view, while at the same time displaying the same conference model data as each of the other client computers.

In the preferred embodiment the server computer does not need to graphically display the conference model data but has a central role in the conference with responsibility for maintaining all the copies of the conference model data on the client computers. Centralised conference model data on the server computer may act as master and force all the client computers to be kept in step.

Hosted machine code is a platform independent code such as the byte code needed for the JAVA Virtual Machine. Hosted machine code such as JAVA byte code provides operating system/platform independence of applications and JAVA byte code is a programmable code that is executable on any JAVA enabled platform (an operating system having a JAVA executable environment or virtual machine) without recompilation. JAVA is a trademark of SUN MICROSYSTEMS CORPORATION.

In order to promote a fuller understanding of this and other aspects of the present invention, an embodiment will now be described, by way of example only, with reference to the accompanying drawings in which:

Figure 1 is prior art example of a peer to peer conferencing system;

Figure 2 is a prior art example of a client/server conferencing system;

Figure 3 is a schematic representation of the embodiment of the invention;

Figure 4 is a schematic representation of the program classes used to implement the embodiment of the invention; and

Figure 5 is a representation of the computer screen for IBM's Teamworker application.

5

10

15

20

A client/server arrangement of the embodiment of the invention has a central server computer (30) connected directly to three client computers (32A, 32B, 32C) (see Figure 3). A client computer (32A) comprises a processor (34) which performs under the instruction of client code (36A, 36B, 36C) (client program code) stored in memory (38) and/or on disk of a client computer. Visual output of the client computer (36A) is through a video peripheral (40) such as a monitor or LCD screen connected to the processor (34) and user input to the client computer (32A) is by a keyboard (42) and computer mouse (43) connected to the processor (34). The client computer further comprises a network adaptor (44) connected to the processor at one end and a computer network (46) at the other. Examples of two computer networks (46) are a corporate intranet and the internet. Examples of the client computer is IBM's Aptiva computer having a 486 133Mhz micro processor, 1Gbyte hard drive, 16Mbytes of memory and a 16 inch monitor.

25

30

The server computer (30) comprises a processor (47) which performs under the instruction of server code (50A, 50B) (server programme code) stored in memory (52) and/or on disk of the server computer. Visual output is not strictly necessary but may be through a video peripheral such as a monitor or LCD screen connected to the processor (48) and user input is also not strictly necessary but would be through a keyboard or computer mouse connected to the processor (46). The server computer (30) further comprises a network adaptor (54) connected to the processor (48) at one end and the computer network (46) at the other. An examples of the client computer is IBM's RISC6000 computer having a Intel P200Mhz micro processor, a 20Gbyte hard drive and 64Mbytes of memory.

35

40

Each instance of client code (36A, 36B, 36C) comprises three main parts: the conference model object (36A) (an instance of a ConferencingAppletModel class), the conference controller object (36B) (an instance of a ConferencingApplet Controller class) and the conference view object (36C) (an instance of a ConferencingAppletView class). The client code forms a client applet (56) which is contained within a World Wide Web page (58) stored on the server. There are as many instances of client code (36) as there are client computers (32A, 32B, 32C) and an

instance of the client code is downloaded and stored on each client computer (32A, 32B, 32C).

5 To download or create an instance of client code on a client computer (32A) a computer having a JAVA enabled web browser connected to the internet is required; on accessing and viewing a conference web page (58) containing a client applet (56) the browser will copy the client applet (56) on to the client computer (32A). The applet (56) will normally execute upon download but may also execute after performing an event such as selecting a button icon displayed on the monitor (40) using the cursor and mouse (43).

10 The server code (50A, 50B) comprises two main parts: the conference manager object (50A) (an instance of a ConferencingSystemManager class) and the conference model object (50B) (an instance of the ConferencingAppletModel class). Other server objects used in the server code are: a ServerMain object (50C); a Connection object (50D); and a ServerMonitor object (50E) (see Figure 4).

20 The ServerMain object (50C) listens for and accepts incoming network connections from clients. When a connection is received it creates an instance of the Connection class (50D). The ServerMain object also creates and owns a single ConferencingSystemManager object.

25 An instance (50B) of the Connection class is created whenever a network connection is accepted by the server so that for each client a Connection object exists. It represents a connection to an individual client. The Connection object (50D) runs a thread which continually listens for incoming messages from the client and passes them onto the ConferencingSystemManager class. The Connection object (50D) is also used for sending messages from the ConferencingSystemManager (50B) back out to the client.

30 The ConferencingSystemManager object (50B) is the object that manages concurrent multi-party conferences for the server. Each conference in progress is represented by an instance of the ConferencingAppletModel class (or to be precise, one of its subclasses). Each conference in progress identifies its model class by a "model id" which is passed on all messages.

The ServerMonitor object (50E) provides a graphical view of the activities of the ConferencingSystemManager (50B). It can be used by an administrator to monitor activity in the conferencing system.

5 The ConferencingAppletModel class (50A) is the superclass for all conference models. Each type of multi-party conference introduces its own subclass of ConferencingAppletModel. For example, there might be a ChatConferenceModel for chat applications, a ChalkboardConferenceModel for chalkboard applications and so on. The model object maintains the
10 current state of the conference in the form of model data (60). It provides three important methods:

 * applyOpCode() (62) - applies an operand and its arguments (as received in a message from a client), to the model, which will cause the model's data (60) (state) to be updated.

15 * setState() - causes the model to set its data (60) (state) to that passed in the message

 * getState() - causes the model to generate a message containing its current state

20 Other client objects used in the client code are: a Sharer object (36D) and a SharerReader object (36E) (see Figure 4).

 The ConferencingAppletModel object (36A) is the same object as used in the server, with subclasses as described above.

25 The ConferencingAppletView class is the superclass for all conference views (used to display the current state of the conference to the user - eg. the appends in a chat conference). Each type of conference introduces its own subclass of ConferencingAppletView. A view
30 class instance (36C) is on one-to-one correspondence with a model. When the model's data (60) changes, it calls the "invalidate()" method of the view which causes the view to redraw itself.

35 The ConferencingAppletController class is the superclass for all conference controllers (used to send messages to and read messages from the server) and forms a controller object (36B). Each type of conference introduces its own subclass of ConferencingAppletController. A
40 controller object (36B) is in one-to-one correspondence with a view and a model. When user actions are initiated at the end-user interface, the view informs the controller of the action to be taken. The controller builds this into a message and sends it to the server (30). The server

(30) then broadcasts the message to all controllers in the conference. Likewise, the controller listens for incoming messages from the server, and passes them to the model's applyOpcode() method (62) for evaluation.

5 The Sharer object (36D) is used by the ConferencingAppletController to establish a socket connection with the server, and then to send messages to it.

10 The SharerReader object (36E) is used by the ConferencingAppletController to listen for incoming messages from the server (30) and pass them to the model for evaluation.

15 The initialisation of the embodiment is as follows. The client computer (32A) under the control of a World Wide Web browser is instructed to view a page (58) containing the client applet (50). The page (58) may be meeting point for users wishing to participate in a certain type of conference. Upon viewing, the client applet (56) is downloaded to the client computer from the server (30) and when executed, instances of the client code are created including a

20 ConferencingAppletModel object (36A), ConferencingAppletController object (36B), ConferencingAppletView object (36C), Sharer object (36D) and SharerReader object (36E). The client computer (32A) under the control of the Sharer object (36D) then attempts to establish a socket connect with the server object.

25 The server code (50A to 50E) for a particular conference is created as soon as a first client requests a connection for a conference that does not exist. When subsequent client computers attempt to establish a socket connection the ServerMain object (50C) creates an instance of the

30 Connection class (50D) and associates it with that client computer. The ServerMain object (50C) determines the state of the model data (60) in the ConferencingAppletModel object (50A) by calling the getState (state) method which creates a message containing the (state). This message is sent by the Connection object (50D) in the server through the network

35 connection to the client object (36A) where it is picked up by the SharerReader object (36E). The message is then passed on to the client's ConferencingAppletModel object (36A) which changes the state of its model data (60) to that sent in the message using its setState (state) method.

40 Once a client computer stores a current model data (60) it receives a constant flow of update messages comprising operation codes (Opcode)

from the server (50A to 50D) requesting the conference model object (36A) to update itself using the applyOpCode(OpCode) method (62). The update messages continue as long as the conference is in progress. The server in fact sends the same update message (OpCode) to each of the clients connected for a particular conference. In this way, if all the conference model objects have the same model data at the start and receive the same message updates (OpCode) in the same sequence then, at any point in the conference, the model data (60) of any conference model object will match.

If a client wishes to partake in the conference by adding or editing the model data (60) then the changes to the model data (60) are expressed as an operation code by the conference controller object (36B) and sent to the conference manager object (50B) in the server computer (30). The conference manager object (50B) then disperses or resends the operation code as an update message to all the partaking client objects (36). Each conference model object (36A) receives the update message and each applyOpCode method (62) updates the each respective set of model data (60).

The conference model object (50A) of the server (30) also receives the update message and the applyOpCode() method (62) update the server object's model data (60).

An embodiment of the invention is IBM's TeamWorker application developed for a corporate intranet as an interactive team working environment. The model data (60) is the data used in a chalkboard application representing a two dimensional surface which may be written or drawn or painted on in many colours. Figure 4 shows the interaction between the main classes in TeamWorker. The other classes in TeamWorker are used to implement particular conferencing applications on top of the TeamWorker infrastructure, and to provide a graphical front end from which those conferences can be launched.

The main applet class (56) that is executed when it is downloaded to a browser. It creates the TeamWorkerWindow used to display the conferencing facilities to the user Figure 5. TeamWorker has a concept of a "team", and each member of the team is represented by an instance of the TeamMemberModel class inside the applet. Every TeamMemberModel has a corresponding TeamMemberView which represents that team member within the main window by displaying their picture in a passport photograph sized

5 window. To the left of the team members photos, a toolbar is displayed which allows the user to launch conferencing facilities. Each conferencing facility available is represented by a button in the toolbar. These buttons are implemented by the classes ChatButton (which is used to start or join a multi-party chat conference, DiaryButton (which is used to display a team members diary), EMailButton (which is used to send e-mail to a team member), HomePageButton (which requests the web browser hosting the team worker applet to display the team members home page), MessageButton (which is used to send a direct message - pops up on the recipients display - to another team member), PhoneButton (used to display another team members phone number), and JoggleButton (which launches a multi-party game called Joggle).

15 Some of these facilities are implemented purely by the client (for example the display of the phone number and home page), and others are true multi-party conferencing applications that require the team worker infrastructure that is the subject of the application.

20 When one of the conference buttons is pressed (for example, the chat button), new ConferenceView (36C), Model (36A) and Controller (36D) instances are created to manage the conference. The actual classes created are subclasses of the ConferencingAppletModel, View and Controller classes that are particular to the given type of conference. In this case we would create instances of the ChatConferenceModel, View and Controller classes. These classes are used to manage the chat in the manner previously described and illustrated in the Figure.

30 The invention is not restricted to chalkboard applications, another embodiment of the invention could be a three dimensional molecular modelling application as used by chemists and biologists and in which the use of the invention would be particularly advantageous. The model data (60) may represent the types, positions and connections of various atoms in a molecule. The conference view object would be able to view the molecule in any orientation and the conference controller object would allow the user to manipulate and change the atoms. Such an application is particularly processor intensive due to the high resolution images use and the model data (60) needs to be quickly accessible to the client computer and storing the model data on the client computer is the main practical solution. This leads to the problems outlined in the introduction to the specification which the application addresses.

40

Due to the nature of programming it is quite possible to implement the invention using a programming language other than an object orientated programming language.

5 In the object orientated environment it is possible for the a client computer to participate in more than one conference.

10 Of course the server computer may in fact be connected to any number of client computers not just three as stated in the example. Further it is not necessary to use the computers given as examples as the invention may be performed using other types of computer connected in different configurations.

15 In summary there is described a personal conferencing method and system for a client/server environment. The server stores conference model data (60) such as a shared chalkboard or a molecular model and each client has a copy of the conference model data (60). When one of said clients edits the model it creates an instruction for operating on the model and sends the instruction to the server. The server operates (62)
20 on its conference model data on receipt of the instruction and resends the instruction to each of the clients and each client performs the same operation (62) on their respective copies of the conference data model (60). Whereby after a plurality of different operating instructions from different clients the respective copies of the conference model data (60)
25 are equivalent.

CLAIMS

1. A personal conferencing method for a server computer (30) having one or more client computers (32A, 32B, 32C) connected comprising the steps of:

(a) storing conference model data (60) on the server computer (30);

(b) operating (62) on the conference model data (60) on receipt of an instruction (64A) from a client computer;

(c) resending the instruction (64B) from the server computer (30) to each of the client computers (32A, 32B, 32C) whereby each client computer (32A, 32B, 32C) performs the same operation (62) on a copy of the conference data model held on the client computer so that all the copies of the conference data model (60) correspond to the conference model data held on the server computer.

2. A personal conferencing method for one or more client computers (32A, 32B, 32C) connected to a server computer (30) comprising the steps of:

(a) each of said client computers storing a copy of conference model data (60);

(b) one of said client computers (32A) creating an instruction (64A) for operating on the conference model data (60) and sending said instruction (64A) to the server computer (30);

(c) each of said client computers (32A, 32B, 32C) receiving a copy of the instruction (64B) from the server computer and operating (62) on their respective models (60) according to the instruction (64B);

whereby after a plurality of different operating instructions (64B) from different clients the respective copies of the conference model data correspond with each other.

3. A method as claimed in claim 1 or 2 wherein the server computer sends the conference model data (60) to a client computer on request.

4. A method as claimed in claim 1,2 or 3 wherein the steps are performed under the control of machine hosted byte code.

5. A method as claimed in claim 4 wherein machine hosted byte code for controlling a client computer (32A, 32B, 32C) is downloadable from the server computer (30) by a network browser running on the client computer (32A, 32B, 32C).

6. A method as claimed in any of claims 2 to 5 wherein said instruction (64A) to operate of the conference model data is in the form of an Opcode with arguments.

7. A personal conferencing system comprising at least two client computers (32A, 32B, 32C) and a server computer (30);

the server computer (36) comprising:

conference model data (60);

means (50A, 62) for receiving an instruction from a client computer and operating on the conference model data according to the instruction;

the client computer (32A, 32B, 32C) comprising:

means (36B) for creating an instruction and sending the instruction to the server computer;

characterised by:

the server computer further comprising means (50B) for dispersing a received instruction to all the client computer (32A, 32B, 32C); and

each of the client computers (32A, 32B, 32C) further comprising:

a copy of the conference model data (60);

means (36A, 62) for receiving a dispersed instruction from the server computer and operating on the copy of the conference model data according to the instruction received from the server computer.

8. A conferencing system as claimed in claim 7 wherein each client computer (32A, 32B, 32C) further comprises means (36C) for viewing the conference model data.

5 9. A personal conferencing system comprising at least two computers connected for message communication; each computer comprising:

conference model data (60);

10 means (36B) for creating an instruction for operating on the conference model data and sending the instruction to another computer;

means (36A, 60) for operating on the conference model data according to an instruction;

15 characterised by:

one of the computers is a server computer (30) and the other computer or computers (32A, 32B, 32C) are client computers, each having a connection to the server computer;

20 the server computer (30) further comprises a means (50B) for dispersing a received instruction to all the connected client computers (32A, 32B, 32C); and

25 the means (36A, 60) for operating comprises means for receiving a dispersed instruction from the server computer;

30 whereby a sequence of operating instructions (64B) from different computers in the personal conferencing system will be performed in the same order on each copy of conference model data (60).

10. A personal conferencing system comprising:

35 a server object and a plurality of client objects;

a model object (36A) having model data (60) and means (62) for operating on the model data (60) to an instruction message (64B);

40 each client object (36) having means (36B) for creating an instruction for operating on the model object (36A);

characterised by:

each client object (36) comprising an associated model object (36A);

5

each client object (36) further comprises means (36B) for passing an operating instruction to the server object;

10

said server object (50) further comprises means (50B) for dispersing a passed instruction from the server object to all client objects (36);

15

thereby allowing a sequence of instructions from different client objects (36) to be performed in the same order by each model object on each instance of the model data (60).

20

11. A personal conferencing system as claimed in claim 10 further comprising a means for packaging a client object and a model object for transmission across a computer network.

25

12. A personal conferencing system as claimed in claim 11 wherein the model object (36A) further comprising means for causing a model object to generate a state message containing the current model data (60) and pass the state message to another object.

30

13. A personal conferencing system as claimed in claim 12 wherein the model object further comprises means for causing a model object to set its model data to that passed in a state message whereby one model object can be copied with another model object.

35

14. A personal conferencing system as claimed in claims 10,11,12 or 13 serving further comprising a centralised model object associated with the server object.

40

15. A personal conferencing system as claimed in claim 14 wherein a model object associated with a client object is copied from the model object associated with the server object.

16. A personal conferencing system as claimed in any of claims 10 to 15 wherein the classes and objects are implemented in a hosted machine byte code.

17. A personal conferencing system as claimed in claim 16 wherein the client class and the model class comprise a hosted machine byte code applet downloadable and executable by a network browser enabled for the hosted machine byte code.



Application No: GB 9707316.7
Claims searched: 1 - 17

Examiner: Paul Nicholls
Date of search: 12 June 1997

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK CI (Ed.O): G4A (APX, AUXX)

Int CI (Ed.6): G06F 9/46, 17/60

Other: Online: WPI

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
X	GB 2,281,423 A (HEWLETT-PACKARD) - See Fig 1	1, 2, 7, 9, 10 at least
X	EP 0,657,833 A2 (IBM) - See Fig 1	1, 2, 7, 9, 10 at least
X	EP 0,645,726 A2 (AT&T) - See Fig 5	1, 2, 7, 9, 10 at least
X	EP 0,645,725 A2 (AT&T) - See Fig 5	1, 2, 7, 9, 10 at least
X	EP 0,558,006 A2 (TOYOTA) - See Fig 1	1, 2, 7, 9, 10 at least
X	EP 0,319,232 A2 (XEROX) - See page 11 lines 37-40	1, 2, 7, 9, 10 at least
X	WO 96/12241 A1 (HAND HELD PRODUCTS) - See page 21 and Fig 4B	1, 2, 7, 9, 10 at least
X	WO 94/11806 A1 (BRIGHAM YOUNG UNIVERSITY) - See page 4	1, 2, 7, 9, 10 at least
X	US 5,392,400 A (BERKOWITZ et al) - Whole document	1, 2, 7, 9, 10 at least

X	Document indicating lack of novelty or inventive step	A.	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.



The Patent Office

23

Application No: GB 9707316.7
Claims searched: 1 - 17

Examiner: Paul Nicholls
Date of search: 12 June 1997

Category	Identity of document and relevant passage	Relevant to claims
X	US 5,319,777 A (PEREZ) - See Fig 12	1, 2, 7, 9, 10 at least
X	US 5,293,619 A (DEAN) - Whole document	1, 2, 7, 9, 10 at least

X Document indicating lack of novelty or inventive step
Y Document indicating lack of inventive step if combined with one or more other documents of same category.
& Member of the same patent family

A Document indicating technological background and/or state of the art.
P Document published on or after the declared priority date but before the filing date of this invention.
E Patent document published on or after, but with priority date earlier than, the filing date of this application.

An Executive Agency of the Department of Trade and Industry

THIS PAGE BLANK (USPTO)